arXiv:1604.00151v2 [cs.SY] 27 Sep 2016

Analysis of gradient descent methods with non-diminishing, bounded errors

Arunselvan Ramaswamy [1] and Shalabh Bhatnagar [2]

[1]arunselvan@csa.iisc.ernet.in
[2]shalabh@csa.iisc.ernet.in
[1,2]Department of Computer Science and Automation, Indian Institute of Science, Bangalore - 560012, India.

September 28, 2016

Abstract

In this paper, we present easily verifiable, sufficient conditions for both stability and convergence (to the minimum set) of gradient descent (GD) algorithms with bounded, non-diminishing errors. These errors often arise from using gradient estimators or because the objective function is noisy to begin with. Our work extends the contributions of Mangasarian & Solodov and Bertsekas & Tsitsiklis. Our framework improves over the aforementioned ones in that both stability (almost sure boundedness) and convergence are guaranteed even in the case of GD with non-diminishing errors. We present a simplified, yet effective implementation of GD using $SPSA$ with constant sensitivity parameters. Further, unlike other papers no 'additional' restrictions are imposed on the step-size and so also on the learning rate when used to implement machine learning algorithms. Finally, we present the results of some experiments to validate the theory.

1 Introduction

Given a continuously differentiable function $f : \mathbb{R}^d \to \mathbb{R}$, we are interested in finding its minimum. The following gradient descent scheme is often employed for this purpose.

$$x_{n+1} = x_n - \gamma(n)\nabla f(x_n), \tag{1}$$

where $\{\gamma(n)\}_{n\geq 0}$ is the given step-size sequence. GD is a popular tool to implement many machine learning algorithms. For example, the backpropagation algorithm for training neural networks employs GD due to its effectiveness and ease of implementation.

When simulating (1), one often uses gradient estimators such as Kiefer-wolfowitz estimator [7], simultaneous perturbation stochastic approximation ($SPSA$) [9], etc., to obtain estimates of the true gradient at each stage which in turn results in *estimation errors* (ϵ_n in (2)). This is particularly true when the form of f or ∇f is unknown. Previously in the literature, convergence of GD with similar errors was studied in [4]. However, their analysis required the

errors to go to zero at the rate of the step-size. Such assumptions are difficult to enforce and may adversely affect the learning rate when employed to implement machine learning algorithms, see Chapter 4.4 of [6]. In this paper, we present sufficient conditions for both stability (almost sure boundedness) and convergence of GD with bounded errors, for which the recursion is given by

$$x_{n+1} = x_n - \gamma_n(\nabla f(x_n) + \epsilon_n). \tag{2}$$

In the above equation, ϵ_n is the estimation error at stage n, further $\|\epsilon_n\| \leq \epsilon$ for all n. Although in traditional GD the errors are deterministic, we do not distinguish between deterministic and stochastic errors. To the best of our knowledge this is the first time an analysis is done for GD with bounded but not necessarily diminishing errors. Further, we do not impose any 'additional' restrictions on the choice of step-size over the standard assumptions, see (A2) in Section 3.1. Our analysis uses techniques developed in the field of viability theory by [1], [2] and [3]. Further, experimental results are presented in Section 5 supporting the theory presented in this paper.

1.1 Our contributions

1. Previous literature such as [4] requires $\|\epsilon_n\| \to 0$ as $n \to \infty$ for it's analysis to work. Further, both [4] and [8] provide conditions that guarantee one of two things (a) GD diverges almost surely or (b) converges to the minimum set almost surely. On the other hand, we only require $\|\epsilon_n\| \leq \epsilon \ \forall \ n$, where $\epsilon > 0$ is fixed a priori. Also, we present conditions under which GD with bounded errors is stable (bounded almost surely) and converges to an arbitrarily small neighborhood of the minimum set almost surely. Note that our analysis works regardless of whether or not $\|\epsilon_n\|$ tends to zero. For more detailed comparisons with [4] and [8] see Section 3.2.

2. Previously, convergence analysis of GD required severe restrictions on the step-size, see [4], [9]. However, in our paper we do not impose any such restrictions on the step-size. See Section 3.2 (specifically points 1 and 3) for more details.

3. Informally, the main result of our paper, Theorem 2, states the following: One wishes to simulate GD with gradient errors that are not guaranteed to vanish over time. As a consequence of allowing non-diminishing errors, one obtains the following: There exists $\epsilon(\delta) > 0$ such that the iterates are stable and converge to the δ-neighborhood of the minimum set (δ being chosen by the simulator) as long as $\|\epsilon_n\| \leq \epsilon(\delta) \ \forall \ n$.

4. In Section 4.2 we discuss how our framework can be exploited to *undertake convenient yet effective implementations of GD*. Specifically, we present an implementation using $SPSA$, although other implementations can be similarly undertaken. In Section 6, we discuss how the results of this paper can be easily extended to a *Newton based implementation of GD*.

2 Definitions used in this paper

[Upper-semicontinuous map] We say that H is upper-semicontinuous, if given sequences $\{x_n\}_{n\geq 1}$ (in $\mathbb{R}^n$) and $\{y_n\}_{n\geq 1}$ (in $\mathbb{R}^m$) with $x_n \to x$, $y_n \to y$

and $y_n \in H(x_n)$, $n \geq 1$, then $y \in H(x)$.
[Marchaud Map] A set-valued map $H : \mathbb{R}^n \to \{subsets\ of\ \mathbb{R}^m\ \}$ is called *Marchaud* if it satisfies the following properties: **(i)** for each $x \in \mathbb{R}^n$, $H(x)$ is convex and compact; **(ii)** *(point-wise boundedness)* for each $x \in \mathbb{R}^n$, $\sup_{w \in H(x)} \|w\| < K\,(1 + \|x\|)$ for some $K > 0$; **(iii)** H is *upper-semicontinuous.*
Let H be a Marchaud map on $\mathbb{R}^d$. The differential inclusion (DI) given by

$$\dot{x} \ \in \ H(x) \tag{3}$$

is guaranteed to have at least one solution that is absolutely continuous. The reader is referred to [1] for more details. We say that $\mathbf{x} \in \sum$ if $\mathbf{x}$ is an absolutely continuous map that satisfies (3). The *set-valued semiflow* Φ associated with (3) is defined on $[0, +\infty) \times \mathbb{R}^d$ as:
$\Phi_t(x) = \{\mathbf{x}(t) \mid \mathbf{x} \in \sum, \mathbf{x}(0) = x\}$. Let $B \times M \subset [0, +\infty) \times \mathbb{R}^d$ and define

$$\Phi_B(M) = \bigcup_{t \in B,\ x \in M} \Phi_t(x).$$

[ω-limit set] Given $M \subseteq \mathbb{R}^d$, the ω-limit *set* is defined as $\omega_\Phi(M) = \bigcap_{t \geq 0} \overline{\Phi_{[t,+\infty)}(M)}$.
[Limit set of a solution] The limit set of a solution $\mathbf{x}$ with $\mathbf{x}(0) = x$ is given by $L(x) = \bigcap_{t \geq 0} \overline{\mathbf{x}([t, +\infty))}$.
[Invariant set] $M \subseteq \mathbb{R}^d$ is *invariant* if for every $x \in M$ there exists a trajectory, $\mathbf{x} \in \sum$, entirely in M with $\mathbf{x}(0) = x$, $\mathbf{x}(t) \in M$, for all $t \geq 0$.
[Open and closed neighborhoods of a set] Let $x \in \mathbb{R}^d$ and $A \subseteq \mathbb{R}^d$, then $d(x, A) := \inf\{\|a - y\| \mid y \in A\}$. We define the *$\delta$-open neighborhood* of A by $N^\delta(A) := \{x \mid d(x, A) < \delta\}$. The *$\delta$-closed neighborhood* of A is defined by $\overline{N^\delta}(A) := \{x \mid d(x, A) \leq \delta\}$. The open ball of radius r around the origin is represented by $B_r(0)$, while the closed ball is represented by $\overline{B}_r(0)$.
[Internally chain transitive set] $M \subset \mathbb{R}^d$ is said to be internally chain transitive if M is compact and for every $x, y \in M$, $\epsilon > 0$ and $T > 0$ we have the following: There exists n and $\Phi^1, \ldots, \Phi^n$ that are n solutions to the differential inclusion $\dot{x}(t) \in h(x(t))$, points $x_1(= x), \ldots, x_{n+1}(= y) \in M$ and n real numbers $t_1, t_2, \ldots, t_n$ greater than T such that: $\Phi^i_{t_i}(x_i) \in N^\epsilon(x_{i+1})$ and $\Phi^i_{[0,t_i]}(x_i) \subset M$ for $1 \leq i \leq n$. The sequence $(x_1(= x), \ldots, x_{n+1}(= y))$ is called an (ϵ, T) chain in M from x to y.
[Attracting set & fundamental neighborhood] $A \subseteq \mathbb{R}^d$ is *attracting* if it is compact and there exists a neighborhood U such that for any $\epsilon > 0$, $\exists\ T(\epsilon) \geq 0$ with $\Phi_{[T(\epsilon),+\infty)}(U) \subset N^\epsilon(A)$. Such a U is called the *fundamental neighborhood* of A. In addition to being compact if the *attracting set* is also invariant then it is called an *attractor*. The *basin of attraction* of A is given by $B(A) = \{x \mid \omega_\Phi(x) \subset A\}$.
[Lyapunov stable] The above set A is Lyapunov stable if for all $\delta > 0$, $\exists\ \epsilon > 0$ such that $\Phi_{[0,+\infty)}(N^\epsilon(A)) \subseteq N^\delta(A)$.
[Upper-limit of a sequence of sets, Limsup] Let $\{K_n\}_{n \geq 1}$ be a sequence of sets in $\mathbb{R}^d$. The *upper-limit* of $\{K_n\}_{n \geq 1}$ is given by, $Limsup_{n \to \infty} K_n := \{y \mid \underline{lim}_{n \to \infty}\, d(y, K_n) = 0\}$.
We may interpret that the lower-limit collects the limit points of $\{K_n\}_{n \geq 1}$ while the upper-limit collects its accumulation points.

3 Assumptions and comparison to previous literature

3.1 Assumptions

Recall that GD with bounded errors is given by the following recursion:

$$x_{n+1} = x_n - \gamma(n)g(x_n), \tag{4}$$

where $g(x_n) \in G(x_n)\ \forall\ n$ and $G(x) := \nabla f(x) + \overline{B}_\epsilon(0)$, $x \in \mathbb{R}^d$. In other words, the gradient estimate at stage n, $g(x_n)$, belongs to an ϵ-ball around the true gradient $\nabla f(x_n)$ at stage n. Note that (4) is consistent with (2) of Section 1. Our assumptions, $(A1)$-$(A4)$ are listed below.

(A1) $G(x) := \nabla f(x) + \overline{B}_\epsilon(0)$ for some fixed $\epsilon > 0$. ∇f is a continuous function such that $\|\nabla f(x)\| \leq K(1 + \|x\|)$ for all $x \in \mathbb{R}^d$, $K > 0$.

(A2) $\{\gamma(n)\}_{n\geq 0}$ is the step-size sequence (learning rate) such that: $\gamma(n) > 0$ $\forall n$, $\sum\limits_{n\geq 0} \gamma(n) = \infty$ and $\sum\limits_{n\geq 0} \gamma(n)^2 < \infty$. Without loss of generality we let $\sup\limits_n \gamma(n) \leq 1$.

Note that G is an upper-semicontinuous map since ∇f is continuous and pointwise bounded. For each $c \geq 1$, we define $G_c(x) := \{y/c \mid y \in G(cx)\}$. Define $G_\infty(x) := \overline{co}\,\langle Limsup_{c\to\infty} G_c(x)\rangle$, see Section 2 for the definition of $Limsup$. Given $S \subseteq \mathbb{R}^d$, the convex closure of S, denoted by $\overline{co}\langle S\rangle$, is the closure of the convex hull of S. It is worth noting that $Limsup_{c\to\infty} G_c(x)$ is non-empty for every $x \in \mathbb{R}^d$. Further, we show that G_∞ is a Marchaud map in Lemma 1. In other words, $\dot{x}(t) \in -G_\infty(x(t))$ has at least one solution that is absolutely continuous, see [1]. Here $-G_\infty(x(t))$ is used to denote the set $\{-g \mid g \in G_\infty(x(t))\}$.

(A3) $\dot{x}(t) \in -G_\infty(x(t))$ has an attractor set $\mathcal{A}$ such that $\mathcal{A} \subseteq B_a(0)$ for some $a > 0$ and $\overline{B}_a(0)$ is a fundamental neighborhood of $\mathcal{A}$.

Since $\mathcal{A} \subseteq B_a(0)$ is compact, we have that $\sup\limits_{x\in\mathcal{A}} \|x\| < a$. Let us fix the following sequence of real numbers: $\sup\limits_{x\in\mathcal{A}} \|x\| = \delta_1 < \delta_2 < \delta_3 < \delta_4 < a$.

(A4) Let $c_n \geq 1$ be an increasing sequence of integers such that $c_n \uparrow \infty$ as $n \to \infty$. Further, let $x_n \to x$ and $y_n \to y$ as $n \to \infty$, such that $y_n \in G_{c_n}(x_n)$, $\forall n$, then $y \in G_\infty(x)$.

It is worth noting that the existence of a global Lyapunov function for $\dot{x}(t) \in -G_\infty(x(t))$ ***is sufficient to guarantee that*** $(A3)$ ***holds. Further,*** $(A4)$ ***is satisfied when*** ∇f ***is Lipschitz continuous.***

Lemma 1. *G_∞ is a Marchaud map.*

Proof. From the definition of G_∞ and G we have that $G_\infty(x)$ is convex, compact and $\sup\limits_{y\in G(x)} \|y\| \leq K(1 + \|x\|)$ for every $x \in \mathbb{R}^d$. It is left to show that G_∞ is an upper-semicontinuous map. Let $x_n \to x$, $y_n \to y$ and $y_n \in G_\infty(x_n)$, for all $n \geq 1$. We need to show that $y \in G_\infty(x)$. We present a proof by contradiction. Since $G_\infty(x)$ is convex and compact, $y \notin G_\infty(x)$ implies that there exists a

linear functional on $\mathbb{R}^d$, say f, such that $\sup_{z \in G_\infty(x)} f(z) \leq \alpha - \epsilon$ and $f(y) \geq \alpha + \epsilon$, for some $\alpha \in \mathbb{R}$ and $\epsilon > 0$. Since $y_n \to y$, there exists $N > 0$ such that for all $n \geq N$, $f(y_n) \geq \alpha + \frac{\epsilon}{2}$. In other words, $G_\infty(x) \cap [f \geq \alpha + \frac{\epsilon}{2}] \neq \phi$ for all $n \geq N$. We use the notation $[f \geq a]$ to denote the set $\{x \mid f(x) \geq a\}$. For the sake of convenience let us denote the set $Limsup_{c \to \infty} G_c(x)$ by $A(x)$, where $x \in \mathbb{R}^d$. We claim that $A(x_n) \cap [f \geq \alpha + \frac{\epsilon}{2}] \neq \phi$ for all $n \geq N$. We prove this claim later, for now we assume that the claim is true and proceed. Pick $z_n \in A(x_n) \cap [f \geq \alpha + \frac{\epsilon}{2}]$ for each $n \geq N$. It can be shown that $\{z_n\}_{n \geq N}$ is norm bounded and hence contains a convergent subsequence, $\{z_{n(k)}\}_{k \geq 1} \subseteq \{z_n\}_{n \geq N}$. Let $\lim_{k \to \infty} z_{n(k)} = z$. Since $z_{n(k)} \in Limsup_{c \to \infty}(G_c(x_{n(k)}))$, $\exists\ c_{n(k)} \in \mathbb{N}$ such that $\|w_{n(k)} - z_{n(k)}\| < \frac{1}{n(k)}$, where $w_{n(k)} \in G_{c_{n(k)}}(x_{n(k)})$. We choose the sequence $\{c_{n(k)}\}_{k \geq 1}$ such that $c_{n(k+1)} > c_{n(k)}$ for each $k \geq 1$.

We have the following: $c_{n(k)} \uparrow \infty$, $x_{n(k)} \to x$, $w_{n(k)} \to z$ and $w_{n(k)} \in G_{c_{n(k)}}(x_{n(k)})$, for all $k \geq 1$. It follows from assumption $(A4)$ that $z \in G_\infty(x)$. Since $z_{n(k)} \to z$ and $f(z_{n(k)}) \geq \alpha + \frac{\epsilon}{2}$ for each $k \geq 1$, we have that $f(z) \geq \alpha + \frac{\epsilon}{2}$. This contradicts the earlier conclusion that $\sup_{z \in h_\infty(x)} f(z) \leq \alpha - \epsilon$.

It remains to prove that $A(x_n) \cap [f \geq \alpha + \frac{\epsilon}{2}] \neq \phi$ for all $n \geq N$. If this were not true, then $\exists \{m(k)\}_{k \geq 1} \subseteq \{n \geq N\}$ such that $A(x_{m(k)}) \subseteq [f < \alpha + \frac{\epsilon}{2}]$ for all k. It follows that $G_\infty(x_{m(k)}) = \overline{co}(A(x_{m(k)})) \subseteq [f \leq \alpha + \frac{\epsilon}{2}]$ for each $k \geq 1$. Since $y_{n(k)} \to y$, $\exists N_1$ such that for all $n(k) \geq N_1$, $f(y_{n(k)}) \geq \alpha + \frac{3\epsilon}{4}$. This is a contradiction. □

3.2 Relevance of our results

(1) Gradient algorithms with errors have been previously studied by Bertsekas and Tsitsiklis [4]. They impose the following restriction on the estimation errors: $\|\epsilon_n\| \leq \gamma(n)(q + p\|\nabla f(x_n)\|)\ \forall\ n$, where $p, q > 0$. If the iterates are stable then $\|\epsilon_n\| \to 0$. In order to satisfy the aforementioned assumption the choice of step-size may be restricted thereby affecting the learning rate (when used within the framework of a learning algorithm). In this paper we analyze the more general and practical case of bounded $\|\epsilon_n\|$ which does not necessarily go to zero. Further *none* of the assumptions used in our paper impose restrictions that affect the step-size.

(2) The main result of Bertsekas and Tsitsiklis [4] states that the GD with errors either diverges almost surely or converges to the minimum set almost surely. An older study by Mangasarian and Solodov [8] shows the exact same result as [4] but for GD without estimation errors ($\epsilon_n = 0\ \forall\ n$). The main results of our paper, Theorems 1 & 2 show that if the GD under consideration satisfies $(A1)$-$(A4)$ then the iterates are stable (bounded almost surely). Further, the algorithm is guaranteed to converge to a *given small neighborhood of the minimum set* provided the estimation errors are bounded by a constant that is a function of the neighborhood size. To summarize, under the more restrictive setting of [4] and [8] the GD is *not* guaranteed to be stable, see the aforementioned references, while the assumptions used in our paper are less restrictive and guarantee stability under the more general setting of bounded error GD. *It may also be noted that ∇f is assumed to be Lipschitz continuous by [4]. This turns out to be sufficient (but not necessary) for $(A1)$ & $(A4)$ to be satisfied.*

(3) The analysis of Spall [9] can be used to analyze a variant of GD that uses

$SPSA$ as the gradient estimator. Spall introduces a sensitivity parameter c_n in order to control the estimation error ϵ_n at stage n. It is assumed that $c_n \to 0$ and $\sum_{n\geq 0}\left(\frac{\gamma(n)}{c_n}\right)^2 < \infty$, see A1, Section III, [9]. Again, this restricts the choice of step-size and affects the learning rate. In this setting our analysis works for the more practical scenario where $c_n = c$ for all n *i.e.,* a constant, see Section 4.2.

(4) The important advancements of this paper are the following: (i) Our framework is more general and practical since the errors are not required to go to zero; (ii) We provide easily verifiable, non-restrictive set of assumptions that ensure almost sure boundedness and convergence of GD and (iii) Our assumptions $(A1)$-$(A4)$ do not affect the choice of step-size.

(5) Our proof technique is of independent interest to the analysis of general recursive inclusions (involving set-valued mean fields) since it is a significant generalization of [5] that only considers regular stochastic approximation.

4 Proof of stability and convergence

We use (4) to construct the linearly interpolated trajectory, $\overline{x}(t)$ for $t \in [0,\infty)$. First, define $t(0) := 0$ and $t(n) := \sum_{i=0}^{n-1}\gamma(i)$ for $n \geq 1$. Then, define $\overline{x}(t(n)) := x_n$ and for $t \in (t(n), t(n+1))$,

$$\begin{aligned}\overline{x}(t) &:= \left(\frac{t(n+1)-t}{t(n+1)-t(n)}\right)\overline{x}(t(n)) \\ &\quad + \left(\frac{t-t(n)}{t(n+1)-t(n)}\right)\overline{x}(t(n+1)).\end{aligned}$$

We also construct the following piece-wise constant trajectory $\overline{g}(t)$, $t \geq 0$ as follows: $\overline{g}(t) := g(x_n)$ for $t \in [t(n), t(n+1))$, $n \geq 0$.

We need to divide time, $[0,\infty)$, into intervals of length T, where $T = T(\delta_2 - \delta_1)+1$. Note that $T(\delta_2-\delta_1)$ is such that $\Phi_t(x_0) \in N^{\delta_2-\delta_1}(\mathcal{A})$ for $t \geq T(\delta_2-\delta_1)$, where $\Phi_t(x_0)$ denotes solution to $\dot{x}(t) \in G_\infty(x(t))$ at time t with initial condition x_0 and $x_0 \in \overline{B}_a(0)$. Note that $T(\delta_2-\delta_1)$ is independent of the initial condtion x_0, see Section 2 for more details. Dividing time is done as follows: define $T_0 := 0$ and $T_n := min\{t(m) : t(m) \geq T_{n-1}+T\}$, $n \geq 1$. Clearly, there exists a subsequence $\{t(m(n))\}_{n\geq 0}$ of $\{t(n)\}_{n\geq 0}$ such that $T_n = t(m(n))$ $\forall\ n \geq 0$. In what follows we use $t(m(n))$ and T_n interchangeably.

To show stability, we use a projective scheme where the iterates are projected periodically, with period T, onto the closed ball of radius a around the origin, $\overline{B}_a(0)$. Here, the radius a is given by $(A3)$. This projective scheme gives rise to the following rescaled trajectories $\hat{x}(\cdot)$ and $\hat{g}(\cdot)$. First, we construct $\hat{x}(t)$, $t \geq 0$: Let $t \in [T_n, T_{n+1})$ for some $n \geq 0$, then $\hat{x}(t) := \frac{\overline{x}(t)}{r(n)}$, where $r(n) = \frac{\|\overline{x}(T_n)\|}{a} \vee 1$ (a is defined in $(A3)$). Also, let $\hat{x}(T_{n+1}^-) := \lim_{t\uparrow T_{n+1}}\hat{x}(t)$, $t \in [T_n, T_{n+1})$. The 'rescaled g iterates' are given by $\hat{g}(t) := \frac{\overline{g}(t)}{r(n)}$.

Let $x^n(t)$, $t \in [0,T]$ be the solution (upto time T) to $\dot{x}^n(t) = -\hat{g}(T_n + t)$, with the initial condition $x^n(0) = \hat{x}(T_n)$, recall the definition of $\hat{g}(\cdot)$ from the beginning of Section 4. Clearly, we have

$$x^n(t) = \hat{x}(T_n) - \int_0^t \hat{g}(T_n+z)\,dz. \tag{5}$$

We begin with a simple lemma which essentially claims that $\{x^n(t), 0 \leq t \leq T \mid n \geq 0\} = \{\hat{x}(T_n + t), 0 \leq t \leq T \mid n \geq 0\}$. The proof is a direct consequence of the definition of $\hat{g}$ and is hence omitted.

Lemma 2. *For all $n \geq 0$, we have $x^n(t) = \hat{x}(T_n + t)$, where $t \in [0, T]$.*

It directly follows from Lemma 2 that $\{x^n(t), t \in [0,T] \mid n \geq 0\} = \{\hat{x}(T_n + t), t \in [0,T] \mid n \geq 0\}$. In other words, the two families of T-length trajectories, $\{x^n(t), t \in [0,T] \mid n \geq 0\}$ and $\{\hat{x}(T_n + t), t \in [0,T] \mid n \geq 0\}$, are really one and the same. When viewed as a subset of $C([0,T], \mathbb{R}^d)$, $\{x^n(t), t \in [0,T] \mid n \geq 0\}$ is equi-continuous and point-wise bounded. Further, from the *Arzela-Ascoli* theorem we conclude that it is relatively compact. In other words, $\{\hat{x}(T_n + t), t \in [0,T] \mid n \geq 0\}$ is relatively compact in $C([0,T], \mathbb{R}^d)$.

Lemma 3. *Let $r(n) \uparrow \infty$, then any limit point of $\{\hat{x}(T_n + t), t \in [0,T] : n \geq 0\}$ is of the form $x(t) = x(0) + \int_0^t g_\infty(s)\, ds$, where $y : [0,T] \to \mathbb{R}^d$ is a measurable function and $g_\infty(t) \in G_\infty(x(t))$, $t \in [0,T]$.*

Proof. For $t \geq 0$, define $[t] := max\{t(k) \mid t(k) \leq t\}$. Observe that for any $t \in [T_n, T_{n+1})$, we have $\hat{g}(t) \in G_{r(n)}(\hat{x}([t]))$ and $\|\hat{g}(t)\| \leq K\,(1 + \|\hat{x}([t])\|)$, since $G_{r(n)}$ is a Marchaud map. Since $\hat{x}(\cdot)$ is the rescaled trajectory obtained by periodically projecting the original iterates onto a compact set, it follows that $\hat{x}(\cdot)$ is bounded a.s. *i.e.,* $\sup_{t \in [0,\infty)} \|\hat{x}(t)\| < \infty$ *a.s.* It now follows from the observation made earlier that $\sup_{t \in [0,\infty)} \|\hat{g}(t)\| < \infty$ *a.s.*

Thus, we may deduce that there exists a sub-sequence of $\mathbb{N}$, say $\{l\} \subseteq \{n\}$, such that $\hat{x}(T_l + \cdot) \to x(\cdot)$ in $C\left([0,T], \mathbb{R}^d\right)$ and $\hat{g}(m(l) + \cdot) \to g_\infty(\cdot)$ weakly in $L_2\left([0,T], \mathbb{R}^d\right)$. From Lemma 2 it follows that $x^l(\cdot) \to x(\cdot)$ in $C\left([0,T], \mathbb{R}^d\right)$. Letting $r(l) \uparrow \infty$ in

$$x^l(t) = x^l(0) - \int_0^t \hat{g}(t(m(l) + z))\, dz,\ t \in [0,T],$$

we get $x(t) = x(0) - \int_0^t g_\infty(z) dz$ for $t \in [0,T]$. Since $\|\hat{x}(T_n)\| \leq 1$ we have $\|x(0)\| \leq 1$.

Since $\hat{g}(T_l + \cdot) \to g_\infty(\cdot)$ weakly in $L_2\left([0,T], \mathbb{R}^d\right)$, there exists $\{l(k)\} \subseteq \{l\}$ such that

$$\frac{1}{N} \sum_{k=1}^{N} \hat{g}(T_{l(k)} + \cdot) \to g_\infty(\cdot) \text{ strongly in } L_2\left([0,T], \mathbb{R}^d\right).$$

Further, there exists $\{N(m)\} \subseteq \{N\}$ such that

$$\frac{1}{N(m)} \sum_{k=1}^{N(m)} \hat{g}(T_{l(k)} + \cdot) \to g_\infty(\cdot) \text{ a.e. on } [0,T].$$

Let us fix $t_0 \in \{t \mid \frac{1}{N(m)} \sum_{k=1}^{N(m)} \hat{g}(T_{l(k)} + t) \to g_\infty(t),\ t \in [0,T]\}$, then

$$\lim_{N(m) \to \infty} \frac{1}{N(m)} \sum_{k=1}^{N(m)} \hat{g}(T_{l(k)} + t_0) = g_\infty(t_0).$$

Since $G_\infty(x(t_0))$ is convex and compact (Proposition 1), to show that $g_\infty(t_0) \in G_\infty(x(t_0))$ it is enough to show $\lim_{l(k)\to\infty} d\left(\hat{g}(T_{l(k)}+t_0), G_\infty(x(t_0))\right) = 0$. Suppose this is not true and $\exists\, \epsilon > 0$ and $\{n(k)\} \subseteq \{l(k)\}$ such that $d\left(\hat{g}(T_{n(k)}+t_0), G_\infty(x(t_0))\right) > \epsilon$. Since $\{\hat{g}(T_{n(k)}+t_0)\}_{k\geq 1}$ is norm bounded, it follows that there is a convergent sub-sequence. For convenience, assume $\lim_{k\to\infty} \hat{g}(T_{n(k)}+t_0) = g_0$, for some $g_0 \in \mathbb{R}^d$. Since $\hat{g}(T_{n(k)}+t_0) \in G_{r(n(k))}(\hat{x}([T_{n(k)}+t_0]))$ and $\lim_{k\to\infty} \hat{x}([T_{n(k)}+t_0]) = x(t_0)$, it follows from assumption $(A4)$ that $g_0 \in G_\infty(x(t_0))$. This leads to a contradiction. □

Note that in the statement of Lemma 3 we can replace '$r(n) \uparrow \infty$' by '$r(k) \uparrow \infty$', where $\{r(k))\}$ is a subsequence of $\{r(n)\}$. Specifically we can conclude that any limit point of $\{\hat{x}(T_k+t), t \in [0,T]\}_{\{k\}\subseteq\{n\}}$ in $C([0,T],\mathbb{R}^d)$, conditioned on $r(k) \uparrow \infty$, is of the form $x(t) = x(0) - \int_0^t g_\infty(z)\, dz$, where $g_\infty(t) \in G_\infty(x(t))$ for $t \in [0,T]$. It should be noted that $g_\infty(\cdot)$ may be sample path dependent (if ϵ_n is stochastic then $g_\infty(\cdot)$ is a random variable). Recall that $\sup_{x\in\mathcal{A}} \|x\| = \delta_1 < \delta_2 < \delta_3 < \delta_4 < a$ (see the sentence following $(A3)$ in Section 3.1). The following technical lemma is an immediate consequence of Lemma 3.

Corollary 1. *$\exists\, 1 < R_0 < \infty$ such that $\forall\, r(l) > R_0$, $\|\hat{x}(T_l+\cdot) - x(\cdot)\| < \delta_3 - \delta_2$, where $\{l\} \subseteq \mathbb{N}$ and $x(\cdot)$ is a solution (up to time T) of $\dot{x}(t) \in -G_\infty(x(t))$ such that $\|x(0)\| \leq 1$. The form of $x(\cdot)$ is as given by Lemma 3.*

Proof. Assume to the contrary that $\exists\, r(l) \uparrow \infty$ such that $\hat{x}(T_l+\cdot)$ is at least $\delta_3 - \delta_2$ away from any solution to the DI. It follows from Lemma 3 that there exists a subsequence of $\{\hat{x}(T_l+t), 0 \leq t \leq T \;:\; l \subseteq \mathbb{N}\}$ guaranteed to converge, in $C([0,T],\mathbb{R}^d)$, to a solution of $\dot{x}(t) \in -G_\infty(x(t))$ such that $\|x(0)\| \leq 1$. This is a contradiction. □

Remark 1. *It is worth noting that R_0 may be sample path dependent. Since $T = T(\delta_2 - \delta_1) + 1$ we get $\|\hat{x}([T_l + T])\| < \delta_3$ for all T_l such that $\|\overline{x}(T_l)\|(= r(l)) > R_0$.*

4.1 Main Results

We are now ready to prove the two main results of this paper. We begin by showing that (4) is stable (bounded a.s.). In other words, we show that $\sup_n \|r(n)\| < \infty$ a.s. Once we show that the iterates are stable we use the main results of Benaïm, Hofbauer and Sorin to conclude that the iterates converge to a closed, connected, internally chain transitive and invariant set of $\dot{x}(t) \in G(x(t))$.

Theorem 1. *Under assumptions $(A1)-(A4)$, the iterates given by (4) are stable i.e., $\sup_n \|x_n\| < \infty$ a.s. Further, they converge to a closed, connected, internally chain transitive and invariant set of $\dot{x}(t) \in G(x(t))$.*

Proof. First, we show that the iterates are stable. To do this we start by assuming the negation *i.e.,* $P(\sup_n r(n) = \infty) > 0$. Clearly, there exists $\{l\} \subseteq \{n\}$ such that $r(l) \uparrow \infty$. Recall that $T_l = t(m(l))$ and that $[T_l + T] = max\{t(k) \mid t(k) \leq T_l + T\}$.

We have $\|x(T)\| < \delta_2$ since $x(\cdot)$ is a solution, up to time T, to the DI given by $\dot{x}(t) \in G_\infty(x(t))$ and $T = T(\delta_2 - \delta_1) + 1$. Since the rescaled trajectory is obtained by projecting onto a compact set, it follows that the trajectory is bounded. In other words, $\sup_{t\geq 0} \|\hat{x}(t)\| \leq K_w < \infty$, where K_w could be sample path dependent. Now, we observe that there exists N such that all of the following happen:
(i) $m(l) \geq N \implies r(l) > R_0$. [since $r(l) \uparrow \infty$]
(ii) $m(l) \geq N \implies \|\hat{x}([T_l + T])\| < \delta_3$. [since $r(l) > R_0$ and Remark 1]
(iii) $n \geq N \implies \gamma(n) < \frac{\delta_4 - \delta_3}{K(1+K_\omega)}$. [since $\gamma(n) \to 0$]

We have $\sup_{x\in\mathcal{A}} \|x\| = \delta_1 < \delta_2 < \delta_3 < \delta_4 < a$ (see the sentence following $(A3)$ in Section 3.1 for more details). Let $m(l) \geq N$ and $T_{l+1} = t(m(l+1)) = t(m(l) + k + 1)$ for some $k > 0$. If $T_l + T \neq T_{l+1}$ then $t(m(l) + k) = [T_l + T]$, else if $T_l + T = T_{l+1}$ then $t(m(l) + k + 1) = [T_l + T]$. We proceed assuming that $T_l + T \neq T_{l+1}$ since the other case can be identically analyzed. Recall that $\hat{x}(T_{n+1}^-) = \lim_{t\uparrow t(m(n+1))} \hat{x}(t)$, $t \in [T_n, T_{n+1})$ and $n \geq 0$. Then,

$$\hat{x}(T_{l+1}^-) = \hat{x}(t(m(l) + k)) - \gamma(m(l) + k)\hat{g}(t(m(l) + k)).$$

Taking norms on both sides we get,

$$\|\hat{x}(T_{l+1}^-)\| \leq \|\hat{x}(t(m(l) + k))\| + \gamma(m(l) + k)\|\hat{g}(t(m(l) + k))\|.$$

As a consequence of the choice of N we get:

$$\|\hat{g}(t(m(l) + k))\| \leq K\,(1 + \|\hat{x}(t(m(l) + k)\|) \leq K\,(1 + K_\omega)\,. \tag{6}$$

Hence,

$$\|\hat{x}(T_{l+1}^-)\| \leq \|\hat{x}(t(m(l) + k))\| + \gamma(m(l) + k)K(1 + K_\omega).$$

In other words, $\|\hat{x}(T_{l+1}^-)\| < \delta_4$. Further,

$$\frac{\|\overline{x}(T_{l+1})\|}{\|\overline{x}(T_l)\|} = \frac{\|\hat{x}(T_{l+1}^-)\|}{\|\hat{x}(T_l)\|} < \frac{\delta_4}{a} < 1. \tag{7}$$

It follows from (7) that $\|\overline{x}(T_{n+1})\| < \frac{\delta_4}{a}\|\overline{x}(T_n)\|$ if $\|\overline{x}(T_n)\| > R_0$. From Corollary 1 and the aforementioned we get that the trajectory falls at an exponential rate till it enters $\overline{B}_{R_0}(0)$. Let $t \leq T_l$, $t \in [T_n, T_{n+1})$ and $n + 1 \leq l$, be the last time that $\overline{x}(t)$ jumps from within $\overline{B}_{R_0}(0)$ to the outside of the ball. It follows that $\|\overline{x}(T_{n+1})\| \geq \|\overline{x}(T_l)\|$. Since $r(l) \uparrow \infty$, $\overline{x}(t)$ would be forced to make larger and larger jumps within an interval of length $T + 1$. This leads to a contradiction since the maximum jump size within any fixed time interval can be bounded using the *Gronwall* inequality. Thus, the iterates are shown to be stable.

It now follows from *Theorem 3.6 & Lemma 3.8* of **Benaïm, Hofbauer and Sorin** [2] that the iterates converge almost surely to a closed, connected, internally chain transitive and invariant set of $\dot{x}(t) \in G(x(t))$. □

Now that the GD with non-diminishing, bounded errors, given by (4), is shown to be stable (bounded a.s.), we proceed to show that these iterates in

fact converge to an arbitrarily small neighborhood of the minimum set. The proof uses *Theorem 2.1* of Benaïm, Hofbauer and Sorin [3] that we state below. Recall that $G(x) = \nabla f(x) + \overline{B}_\epsilon(0)$, see $(A3)$ of Section 3.1. Let the minimum set, $\mathcal{M}$, of f be the global attractor of $\dot{x}(t) = -\nabla f(x(t))$. It can be shown that any compact set, $\mathcal{M} \subseteq \mathcal{K} \subset \mathbb{R}^d$, is a fundamental neighborhood of $\mathcal{M}$ (see Section 2 for the definition of fundamental neighborhood). It follows from Theorem 1 that the iterates are bounded almost surely. In other words, $\overline{x}(t) \in \mathcal{K}_0$, $\forall\ t \geq 0$, for some compact set $\mathcal{K}_0$ that could be sample path dependent. Hence, $\mathcal{K}_0$ is a fundamental neighborhood of $\mathcal{M}$. Suppose f has several local minima then one needs to consider the (local) minimum set whose fundamental neighborhood is $\mathcal{K}_0$ instead of $\mathcal{M}$. We are now ready to present *Theorem 2.1*, [3]. The statement has been interpreted to the setting of this paper for the sake of convenience.

[***Theorem 2.1***, [3]] Given $\delta > 0$, there exists $\epsilon(\delta) > 0$ such that there exists a unique attractor $\mathcal{M}'$ of the DI $\dot{x}(t) \in -\left(\nabla f(x(t)) + \overline{B}_r(0)\right)$ with $\mathcal{M}' \subseteq N^\delta(\mathcal{M})$ provided $\nabla f(x) + \overline{B}_r(0) \subseteq N^{\epsilon(\delta)}(\nabla f(x))$ for each $x \in \mathbb{R}^d$, where $r \geq 0$. Further, $\mathcal{K}_0$ is also the fundamental neighborhood associated with $\mathcal{M}'$.

Theorem 2. *Given $\delta > 0$, there exists $\epsilon(\delta) > 0$ such that the GD with bounded errors given by (4) converges to $N^\delta(\mathcal{M})$, the δ-neighborhood of the minimum set of f, provided $\epsilon \leq \epsilon(\delta)/2$. Here ϵ is the bound for estimation errors from assumption $(A1)$.*

Proof. As stated earlier we have the following: (a) $\overline{x}(t) \in \mathcal{K}_0$, $\forall\ t \geq 0$ and (b) The minimum set of f, $\mathcal{M}$, is the global attractor of $\dot{x}(t) = -\nabla f(x(t))$ such that $\mathcal{K}_0$ is its fundamental neighborhood. It follows from *Theorem 2.1* of [3] that there exists $\epsilon(\delta) > 0$ such that $\dot{x}(t) \in -\left(\nabla f(x(t)) + \overline{B}_r(0)\right)$ has an attractor $\mathcal{M}' \subseteq N^\delta(\mathcal{M})$ with fundamental neighborhood $\mathcal{K}_0$ provided $r < \epsilon(\delta)$. Let us fix $\epsilon := \epsilon(\delta)/2$ in $G = \nabla f + \overline{B}_\epsilon(0)$, see $(A1)$ in Section 3.1. Since $\epsilon < \epsilon(\delta)$, the DI $\dot{x}(t) \in -G(x(t))$ has an attractor $\mathcal{M}_1 \subseteq N^\delta(\mathcal{M})$ with $\mathcal{K}_0$ as the fundamental neighborhood. We know that the iterates given by (4) track a solution to $\dot{x}(t) \in G(x(t))$, see *Proposition 1.3* of [2]. In other words, the iterates converge to $\mathcal{M}_1$ since it is the attractor of $\dot{x}(t) \in G(x(t))$. Further, $\mathcal{M}_1 \subseteq N^\delta(\mathcal{M})$ *i.e.,* the iterates converge to the δ-neighborhood of the minimum set, $\mathcal{M}$. □

4.2 Implementing GD methods using SPSA

Gradient estimators are often used in the implementation of GD methods, both for convenience and ease of implementation. In this section, we consider an implementation of GD using $SPSA$, [9]. When using $SPSA$ the update rule for the i^{th} coordinate is given by

$$x_{n+1}^i = x_n^i - \gamma(n)\left(\frac{f(x_n + c_n\Delta_n) - f(x_n - c_n\Delta_n)}{2c_n\Delta_n^i}\right), \tag{8}$$

where $x_n = \left(x_n^1, \ldots, x_n^d\right)$ is the underlying parameter, $\Delta_n = \left(\Delta_n^1, \ldots, \Delta_n^d\right)$ is a sequence of perturbation random vectors such that Δ_n^i, $1 \leq i \leq d$, $n \geq 0$ are *i.i.d.*. It is common to assume Δ_n^i to be symmetric, Bernoulli distributed, taking values ± 1 *w.p.* $1/2$. The sensitivity parameter c_n is such that the following are assumed: $c_n \to 0$ as $n \to \infty$; $\sum_{n\geq 0}\left(\frac{\gamma(n)}{c_n}\right)^2 < \infty$, see $A1$ of [9]. Further, c_n

needs to be chosen such that the estimation errors go to zero. This, in particular, could be difficult since the form of the function f is often unknown. One may need to run experiments to find each c_n. Also, smaller values of c_n in the initial iterates tend to blow up the variance which in turn affects convergence. For these aforementioned reasons, in practice, one lets $c_n := c$ (a small constant) for all n. If we assume additionally that the second derivative of f is bounded, then it is easy to see that the estimation errors are bounded by $\epsilon(c)$ such that $\epsilon(c) \to 0$ as $c \to 0$. ***Thus, it is clear that keeping*** c_n ***fixed to*** c ***forces the estimation errors to be bounded at each stage. In other words, SPSA with a constant sensitivity parameter falls under the purview of the framework presented in this paper.*** Also, it is worth noting that the iterates are assumed to be stable (bounded a.s.) in [9]. However in our framework, stability is shown under verifiable conditions even when $c_n = c,\ n \geq 0$.

We arrive at the important question of how to choose this constant c in practice such that fixing $c_n := c$ we still get the following: (a) the iterates are stable and (b) GD implemented in this manner converges to a minimum. Suppose the simulator wants to ensure that the iterates converge to a δ-neighborhood of the minimum *i.e.,* $N^\delta(\mathcal{M})$, then it follows from Theorem 2 that there exists $\epsilon(\delta) > 0$ such that the GD converges to $N^\delta(\mathcal{M})$ provided the estimation error at each stage is bounded by $\epsilon(\delta)$. Now, c is chosen such that $\epsilon(c) \leq \epsilon(\delta)$. The simulation is carried out by fixing the sensitivity parameters to this c. As stated earlier one may need to carry out experiments to find such a c. However, the advantage is that we only need to do this once before starting the simulation. Also, the iterates are guaranteed to be stable and converge to the δ-neighborhood of the minimum set provided $(A1)$-$(A4)$ are satisfied.

5 Experimental results

In this section, we present the results of two experiments to support the theory presented in this paper. For purposes of illustration and simplicity we consider the quadratic objective function $f : \mathbb{R}^d \to \mathbb{R}$ with $f(x) := x^T Q x$, where Q is a positive definite matrix. Clearly, the origin is the unique global minimizer of f. The minimum set of f can be found using the following GD scheme. This scheme uses $SPSA$ with constant sensitivity parameter c for gradient estimations.

$$x_{n+1} = x_n - a(n) \begin{bmatrix} \frac{f(x_n + c\Delta_n) - f(x_n - c\Delta_n)}{2c\Delta_{n,1}} \\ \vdots \\ \frac{f(x_n + c\Delta_n) - f(x_n - c\Delta_n)}{2c\Delta_{n,d}} \end{bmatrix}. \tag{9}$$

In the above, $\{\Delta_{n,i} \mid n \geq 0,\ 1 \leq i \leq d\}$ is a sequence of *i.i.d.* symmetric Bernoulli (± 1 *w.p.* $1/2$) random variables and $\{a(n)\}_{n \geq 0}$ is the chosen step-size sequence.

Parameter settings:

1. The dimension $d = 10$. The positive definite matrix Q was randomly generated.

2. The number of iterations of (9) was 8000. The starting point x_0 was randomly chosen.

3. c was varied from 0.1 to 10. For each value of c, the recursion given by (9) was run for 8000 iterations and $\|x_{8000}\|$ was recorded. Since origin is the unique global minimizer of f, $\|x_{8000}\|$ records the distance of the iterate after 8000 iterations from the origin.

4. For $0 \leq n \leq 7999$, we chose the following step-size sequence: $a(n) = \frac{1}{(n \mod 800)+100}$, $n \geq 1$.

The above step-size sequence was chosen since it seems to expedite the convergence of the iterates to the minimum set. We were able to use this sequence since our framework does not impose any restrictions on step-sizes. The freedom to choose step-sizes illustrates one of the main advantages of our framework over previous ones, see [9]. Further, since we keep the sensitivity parameters fixed the implementation was greatly simplified. Based on the theory presented in this paper, for larger values of c one expects the iterates to be farther from the origin than for smaller values of c. This theory is corroborated by the experiment illustrated in Fig. 1. Note that to generate Q we first generate a column-orthonormal matrix U and let $Q := U\Sigma U^T$, where Σ is a diagonal matrix with strictly positive entries. To generate U we sample it's entires independently from a Gaussian distribution then we apply Gram-Schmidt orthogonalization to the columns.

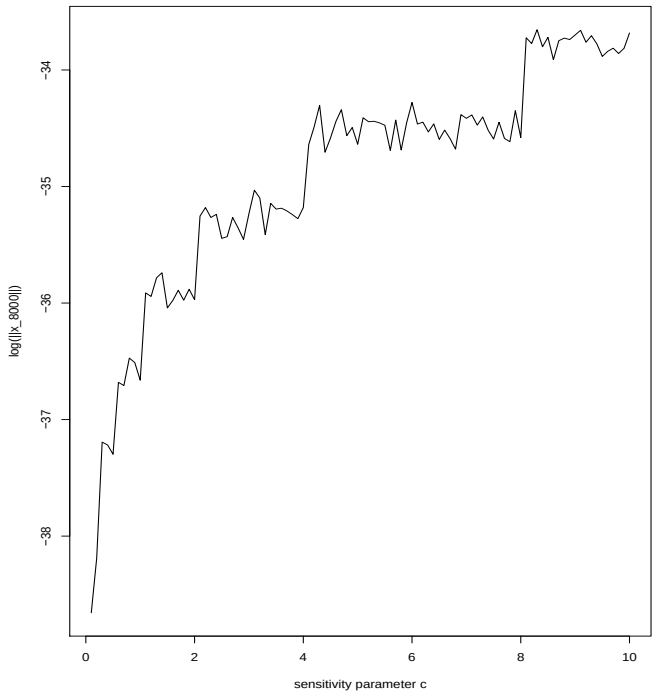

Figure 1: Performance variation as a function of the sensitivity parameter c.

In Fig. 1 the x-axis represents the values of c ranging from 0.1 to 10 while the y-axis represents *logarithm of the corresponding distances from the origin after* 8000 *iterations i.e.,* $\log(\|x_{8000}\|)$. Note that for c close to 0, the iterate,

x_{8000} belongs to the $B_{e^{-38}}(0)$ neighborhood of the origin while for c close to 10, the iterate, x_{8000} only belongs to the $B_{e^{-32}}(0)$ neighborhood. Also note that the graph has a series of "steep rises" followed by "plateaus". These indicate that for values of c within the same plateau the iterate converges to the same neighborhood of the origin. As stated earlier for larger values of c the iterates are farther from the origin than for smaller values of c.

For the second experiment we ran the following recursion for 1000 iterations:

$$x_{n+1} = x_n + 1/n\,(Qx_n + \underline{\epsilon})\,, \text{ where} \tag{10}$$

1. the starting point x_0 was randomly chosen. The dimension $d = 10$.
2. the matrix Q was a randomly generated positive definite matrix (Q is generated as explained before),
3. $\underline{\epsilon} = \begin{bmatrix} \epsilon/\sqrt{d} \\ \vdots \\ \epsilon/\sqrt{d} \end{bmatrix}$, is the constant noise-vector added at each stage and $\epsilon \in \mathbb{R}$.

Since Q is a positive definite matrix, we expect the recursion given by (10) to converge to the origin when $\epsilon = 0$ in the noise-vector. A natural question to ask is the following: If a "small" noise-vector is added at each stage does the iterate still converge to a small neighborhood of the origin or do the iterates diverge? It can be verified that (10) satisfies $(A1)$-$(A4)$ of Section 3.1 for any $\epsilon \in \mathbb{R}$. Hence it follows from Theorem 1 that the iterates are stable and do not diverge. In other words, the addition of such a noise does not accumulate and force the iterates to diverge. As in the first experiment we expect the iterates to be farther from the origin for larger values of ϵ. This is evidenced by the plot in Fig. 2.

The x-axis in fig. 2 represents values of the ϵ parameter in (10) that varies from 0.1 to 2 *i.e.,* $\|\underline{\epsilon}\|$ varies from 0.1 to 2. The y-axis represents the distance of the iterate from the origin after 1000 iterations *i.e.,* $\|x_{1000}\|$. For ϵ close to 0 the iterate (after 1000 iterations) is within $B_{0.0003}(0)$ while for ϵ close to 2 the iterate (after 1000 iterations) is only within $B_{0.1}(0)$.

6 Extensions and conclusions

The main results of this paper can be extended to analyze another popular implementation of GD using Newton's method with non-diminishing, bounded errors. To see this define $G(x) := H(x)^{-1}\nabla f(x) + \overline{B}_\epsilon(0)$ in $(A1)$; G_∞ changes accordingly. Here $H(x)$ (assumed positive definite) denotes the Hessian evaluated at x. Theorems 1 & 2 hold under this new definition of G and appropriate modifications of $(A1) - (A4)$. Another extension to our main results is the introduction of an additional martingale noise term M_{n+1} at stage n. The main results of this paper will continue to hold provided $\sum_{n\geq 0}\gamma(n)M_{n+1} < \infty$ a.s. To summarize, in this paper we provide sufficient conditions for stability and convergence of GD with non-diminishing, bounded errors. To the best of our knowledge this is the first time GD with bounded errors has been analyzed. In addition to being easily verifiable, the assumptions presented herein do not affect the choice of step-size. Finally, experimental results are seen to validate the theory presented in Section 5.

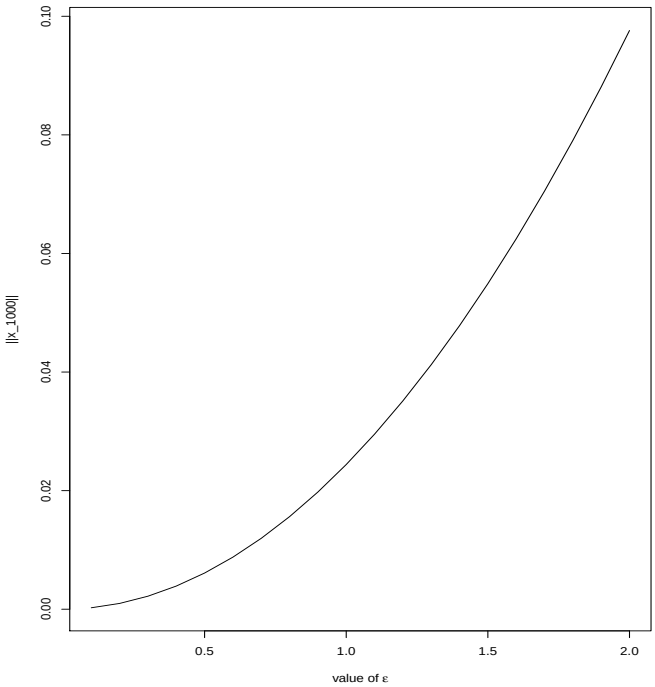

Figure 2: Variation in performance as a function of the neighborhood parameter ϵ.

References

[1] J. Aubin and A. Cellina. *Differential Inclusions: Set-Valued Maps and Viability Theory*. Springer, 1984.

[2] M. Benaïm, J. Hofbauer, and S. Sorin. Stochastic approximations and differential inclusions. *SIAM Journal on Control and Optimization*, pages 328–348, 2005.

[3] M. Benaïm, J. Hofbauer, and S. Sorin. Perturbations of set-valued dynamical systems, with applications to game theory. *Dynamic Games and Applications*, 2(2):195–205, 2012.

[4] D.P. Bertsekas and J.N. Tsitsiklis. Gradient convergence in gradient methods with errors. *SIAM Journal on Optimization*, 10(3):627–642, 2000.

[5] V. S. Borkar and S.P. Meyn. The O.D.E. method for convergence of stochastic approximation and reinforcement learning. *SIAM J. Control Optim*, 38:447–469, 1999.

[6] S.S. Haykin. *Neural networks and learning machines*, volume 3. Pearson Education Upper Saddle River, 2009.

[7] J. Kiefer and J. Wolfowitz. Stochastic estimation of the maximum of a regression function. *The Annals of Mathematical Statistics*, 23(3):462–466, 1952.

[8] O.L. Mangasarian and M.V. Solodov. Serial and parallel backpropagation convergence via nonmonotone perturbed minimization. *Optimization Methods and Software*, 4(2):103–116, 1994.

[9] J.C. Spall. Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *Automatic Control, IEEE Transactions on*, 37(3):332–341, 1992.